

# MULTIGROUP MODELS OF THE CONVECTIVE EPOCH IN CORE COLLAPSE SUPERNOVAE \*

F. DOUGLAS SWESTY AND ERIC S. MYRA

*Dept. of Physics & Astronomy*  
*State University of New York at Stony Brook*  
*Stony Brook, NY 11794-3800 USA*  
*E-mail: dswesty@mail.astro.sunysb.edu*  
*emyra@mail.astro.sunysb.edu*

Understanding the explosion mechanism of core collapse supernovae is a problem that has plagued nuclear astrophysicists since the first computational models of this phenomenon were carried out in the 1960s. Our current theories of this violent phenomenon center around multi-dimensional effects involving radiation-hydrodynamic flows of hot, dense matter and neutrinos. Modeling these multi-dimensional radiative flows presents a computational challenge that will continue to stress high-performance computing beyond the teraflop to the petaflop level. In this paper we describe a few of the scientific discoveries that we have made via terascale computational simulations of supernovae under the auspices of the SciDAC-funded Terascale Supernova Initiative.

## 1. Introduction: Scientific and Computational Challenges

For over four decades, researchers have struggled with one of the great unsolved problems of astrophysics: how massive stars end their lives via collapse of their cores and generate the explosions we observe as supernovae.

The problem is an important one for diverse branches of physics. The supernova process is ultimately responsible for the distribution of nearly all elements in the universe heavier than carbon. Supernovae frequently leave a compact remnant whose densities are at or above nuclear densities. The majority of energy released during gravitational collapse is imparted to weakly interacting neutrinos. Finally, a collapse departing from spherical symmetry can generate gravitational waves. Therefore, supernovae serve as natural laboratories for matter at the highest densities and sit at the frontier of research in astronomy and nuclear, particle, and gravitational physics.

The supernova problem can be briefly stated. Through nuclear fusion, stars more massive than about 10 solar masses ( $M_{\odot}$ ) evolve to where the center consists of an iron core surrounded by layers of successively lighter elements. There is no exothermic reaction that can process iron, and the core becomes dynamically unstable as a result of electron capture by protons. Collapse ensues, but decelerates once nuclear density is exceeded because of the strong repulsive force of nucleons. The core bounces and rebounds, producing a shock wave that starts propagating outwards through the mantle. Were the shock to continue moving in this fashion, matter would be explosively ejected, and the supernova problem would be solved. Instead, however, the most realistic supernova models collapse and rebound, but create a shock wave that doesn't eject matter, either on the hydrodynamic

---

\*Invited talk, presented at SciDAC 2005, San Francisco, CA, USA, 26–30 June 2005; to appear in *Journal of Physics: Conference Series*.

timescale ( $\sim 10$  ms) or the diffusive timescale of the escaping neutrino radiation ( $\sim 1\text{--}10$  s).

This longstanding supernova puzzle has a longstanding history of challenging the highest-performance computer systems of each technology generation. The difficulty of assembling a convincing supernova model has led researchers to add increasingly sophisticated and computationally intensive physics to the mix. Yet, despite great advances in the field, supernova models continue to be oversimplified and leave many relevant branches of physics inadequately explored. The SciDAC program has enabled physicists to make substantial progress in closing this gap and has permitted both fundamental research as well as software development, allowing the field to progress at a more optimistic pace.

In this article, we present recent results of core-collapse simulations performed as part of the Terascale Supernova Initiative under SciDAC. These simulations were performed using 1024 IBM POWER-3 processors on *seaborg*, the IBM SP at the National Energy Research Scientific Computing Center (NERSC). Typical simulations consumed between 50,000 and 100,000 processor-hours to advance 30 ms of model time.

These are the first supernova simulations performed in a two-dimensional geometry that also include a detailed multigroup treatment of neutrino transport. Our current focus is the 30 ms immediately following core bounce. This epoch gives rise to dynamic instabilities in areas behind the stalled shock and leads to complex convective behavior of the stellar fluid. This is also a time when neutrino radiation plays a major role in determining the all-important transport of energy within the collapsed core.

## 2. Our Model for a Core-Collapse Supernova Simulation

Our model includes a radiation-hydrodynamic simulation algorithm to solve the equations of (1) hydrodynamic and (2) neutrino-radiation evolution. The model also requires input of (3) initial conditions, including a pre-collapse iron stellar core, together with several of the surrounding layers of lighter elements; (4) an equation of state to describe the behavior of nuclear matter over the diverse density, temperature, and composition regimes that are encountered; and (5) neutrino microphysics to describe the weak interactions of expected importance.

The computer code that implements this model, V2D, is a two-dimensional, pure Eulerian, staggered-mesh code, loosely based on the ZEUS-2D algorithm by Stone and Norman [7,8,9], but considerably extended to permit its more general use in radiation hydrodynamic problems.

V2D, an entirely new implementation coded according to the Fortran 95 standard, includes the design goals of portability, scalability, componentization, and adherence to standards. Designed for distributed-memory parallel architectures, V2D uses MPI-1 for message passing between application tasks. The SciDAC program has largely enabled the development of V2D's high degree of scalability. The range of system sizes on which it currently runs includes standalone Linux-based laptops and 2048 processors of *seaborg*, the IBM POWER-3-based SP currently installed at NERSC. V2D formats data input and output using parallel HDF5, built on the MPI-I/O portion of the MPI-2 standard.

Consistent with our goal of componentization, we insist on a complete separation of microphysics from the numerical implementation of our radiation-hydrodynamics algorithm. This isolation of physics components from the mathematics and computational-science components has resulted in significant algorithmic and code improvements contributed by our non-physicist colleagues.

## 2.1. Hydrodynamics

V2D’s hydrodynamic component is solved in Newtonian formalism, wherein the Eulerian equations are explicitly differenced. For core collapse, the problem is posed in spherical polar geometry and expressed in two spatial dimensions by assuming azimuthal symmetry. To accomplish this, we use the following set of hydrodynamic equations:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \quad (1)$$

$$\frac{\partial (\rho Y_e)}{\partial t} + \nabla \cdot (\rho Y_e \mathbf{v}) = -m_b \sum_f \int d\varepsilon \left( \frac{\mathbb{S}_\varepsilon}{\varepsilon} - \frac{\bar{\mathbb{S}}_\varepsilon}{\varepsilon} \right) \quad (2)$$

$$\frac{\partial E}{\partial t} + \nabla \cdot (E \mathbf{v}) + P \nabla \cdot \mathbf{v} = - \sum_f \int d\varepsilon (\mathbb{S}_\varepsilon + \bar{\mathbb{S}}_\varepsilon) \quad (3)$$

$$\frac{\partial (\rho \mathbf{v})}{\partial t} + \nabla \cdot (\rho \mathbf{v} \mathbf{v}) + \nabla P + \rho \nabla \Phi + \nabla \cdot \left\{ \sum_f \int d\varepsilon (P_\varepsilon + \bar{P}_\varepsilon) \right\} = 0. \quad (4)$$

Equation (1) is the continuity equation for mass, where  $\rho$  is the mass density and  $\mathbf{v}$  is the matter velocity, and where these quantities, and those in the following equations, are understood to be functions of position  $\mathbf{x}$  and time  $t$ . Equation (2) expresses the evolution of electric charge, where  $Y_e$  is the ratio of the net number of electrons over positrons to the total number of baryons. In the presence of weak interactions, the right-hand side is non-zero to account for reactions where the number of electrons can change. Here, we express the net emissivity of a neutrino flavor (of energy  $\varepsilon$ ) and its antineutrino by  $\mathbb{S}_\varepsilon$  and  $\bar{\mathbb{S}}_\varepsilon$ , respectively. This expression is integrated over all neutrino energies and summed over all neutrino flavors  $f$ . The mean baryonic mass is given by  $m_b$ . Evolution of the internal energy of the matter is given by the gas-energy equation, Eq. (3), where  $E$  is the matter internal energy density and  $P$  is the matter pressure. Again, the right-hand side of this equation is non-zero whenever energy is transferred between matter and neutrino radiation as a result of weak interactions. Finally, Eq. (4) expresses gas-momentum conservation, where  $\Phi$  is the gravitational potential, and  $P_\varepsilon$  and  $\bar{P}_\varepsilon$  are radiation-pressure tensors for each energy and flavor of neutrino and its anti-neutrino, respectively.

The combined use of explicit hydrodynamics, spherical polar coordinates, and a spatial domain that includes the origin presents a numerical stability problem. This follows from the degeneracy of the coordinate system at  $r = 0$ . Clearly, all central zones at the vertex should be in mutual instantaneous sonic contact. However, since standard explicit nearest-neighbor finite-difference techniques do not include the coupling of all zones at the vertex, we introduce numerical “baffles” into the center of the collapsed core, as though it were a tank of fluid. These baffles prevent both the angular motion of the fluid inside the baffle radius and the severe timestep-restricting Courant-Friedrichs-Levy (CFL) condition in the  $\theta$  direction near the core’s center as a result of the converging zones. Zones on either side of a baffle remain sonically connected, but only in that sound waves flow around the outer edge of that baffle. Baffles are kept small to prevent interference with any developing proto-neutron star (PNS) instability. In current models, we place the outer radius of the baffle in the range of about 4–8 km, which permits a timestep of about  $2\text{--}5 \times 10^{-7}$  s.

## 2.2. Neutrino Transport

Because supernova neutrinos, in general, cannot be described by an equilibrium distribution function, neutrino transport is the most difficult component of a supernova model to implement. A solution requires a complete phase-space description of each neutrino's position and momentum and, thus, a solution of the six-dimensional Boltzmann transport equation or some reasonable approximation of it. This high dimensionality causes the transport portion of a model to have the largest share of computational cost in execution time, computer memory, and I/O.

The Boltzmann transport equation can be expressed in terms of the radiation intensity,  $I = I(\varepsilon, \mathbf{x}, \mathbf{\Omega}, t)$ , where  $\varepsilon$  is the energy of a neutrino,  $\mathbf{x}$  its position, and  $\mathbf{\Omega}$  the solid angle into which the neutrino radiation is directed. In terms of  $I$ , the Newtonian form can be expressed as

$$\frac{1}{c} \frac{\partial I}{\partial t} + \mathbf{\Omega} \cdot \nabla I + \sum_i a_i \frac{\partial I}{\partial p_i} = \left( \frac{\partial f}{\partial t} \right)_{\text{coll.}}, \quad (5)$$

where  $a_i$  is the  $i^{\text{th}}$  component of the matter acceleration,  $p_i$  the  $i^{\text{th}}$  component of the momentum of the neutrino, and  $c$  the speed of light. The right-hand side of Eq. (5) lumps together the contributions from all interactions that a neutrino might experience and is collectively referred to as the collision integral.

In V2D, we implement a fully two-dimensional, multi-group flux-limited diffusion scheme. This scheme extends our earlier work [6, 11] and is accomplished by taking the zeroth angular moment of the Eq. (5) to yield the following neutrino monochromatic energy equation in the co-moving frame

$$\frac{\partial E_\varepsilon}{\partial t} + \nabla \cdot (E_\varepsilon \mathbf{v}) + \nabla \cdot \mathbf{F}_\varepsilon - \varepsilon \frac{\partial}{\partial \varepsilon} (P_\varepsilon : \nabla \mathbf{v}) = \mathbb{S}_\varepsilon, \quad (6)$$

where  $E_\varepsilon$  is the zeroth angular moment of  $I$  and is also the neutrino energy density per unit energy interval at position  $\mathbf{x}$  and time  $t$ . The first angular moment of  $I$  is given by  $\mathbf{F}_\varepsilon$ , which is also the neutrino energy flux per unit energy interval. The quantities  $P_\varepsilon$  and  $\mathbb{S}_\varepsilon$  are as previously defined ( $P_\varepsilon$  being additionally the second angular moment of  $I$ ). The expression  $P_\varepsilon : \nabla \mathbf{v}$  indicates contraction in both indices of the second-rank tensors  $P_\varepsilon$  and  $\nabla \mathbf{v}$ . A corresponding equation describes the antineutrinos, which can be expressed by substituting the antineutrino energy density  $\bar{E}_\varepsilon$ , and corresponding higher moments, for each instance of  $E_\varepsilon$  and *vice versa* in Eq. (6). This pair of equations is solved for each neutrino energy  $\varepsilon$  and neutrino flavor over the entire computational domain. We currently track electronic, muonic, and tauonic neutrinos.

The right-hand side of Eq. (6) can be expressed as

$$\begin{aligned} \mathbb{S}_\varepsilon \equiv & S_\varepsilon \left( 1 - \frac{\alpha}{\varepsilon^3} E_\varepsilon \right) - c \kappa_\varepsilon^a E_\varepsilon + \left( 1 - \frac{\alpha}{\varepsilon^3} E_\varepsilon \right) \varepsilon \int d\varepsilon' G(\varepsilon, \varepsilon') \left( 1 - \frac{\alpha}{\varepsilon'^3} \bar{E}_{\varepsilon'} \right) \\ & + \left( 1 - \frac{\alpha}{\varepsilon^3} E_\varepsilon \right) c \int d\varepsilon' \kappa^{\text{in}}(\varepsilon, \varepsilon') E_{\varepsilon'} - E_\varepsilon c \int d\varepsilon' \kappa^{\text{out}}(\varepsilon, \varepsilon') \left( 1 - \frac{\alpha}{\varepsilon'^3} E_{\varepsilon'} \right), \end{aligned} \quad (7)$$

where  $S_\varepsilon$  and  $\kappa_\varepsilon^a$  are the neutrino emissivities and absorption opacities, respectively, for electron capture processes—a reaction that is only significant for electron neutrinos. Production of neutrino-antineutrino pairs is accounted for in the term containing the pair production kernel  $G(\varepsilon, \varepsilon')$ . Finally, the last pair of terms accounts for non-conservative scattering, processes that can scatter neutrinos into an energy state  $\varepsilon$  from energy state  $\varepsilon'$  with opacity  $\kappa^{\text{in}}(\varepsilon, \varepsilon')$  and the reverse process with opacity  $\kappa^{\text{out}}(\varepsilon, \varepsilon')$ . Once again, the analogous expressions for antineutrinos can be expressed with the substitutions noted above. The phase-space factor  $\alpha$  equals  $(hc)^3/4\pi$  for neutrinos, where  $h$  is Planck's constant.

Equation (6) is closed using Levermore and Pomraning's prescription for flux-limited diffusion [5], which allows us to express  $\mathbf{F}_\varepsilon = -D_\varepsilon \nabla E_\varepsilon$ , where  $D_\varepsilon$  is a “variable” diffusion coefficient

that yields the correct fluxes for the diffusion and free-streaming limits and an approximate solution in the intermediate regime. This prescription also provides the elements of the radiation-pressure tensor  $P_\epsilon$ , which closes the system.

Since the neutrino CFL condition is far too restrictive to permit an explicit transport solution, we use a purely implicit method. The equations comprising each neutrino-antineutrino species are assembled into linear-system form. Blocking terms arising from Fermi-Dirac statistical restrictions on final neutrino states—the  $1 - \alpha E_\epsilon/\epsilon^3$  terms in Eq. (7)—make this a system of non-linear equations. Fortunately, the system is sparse, making it amenable to an iterative solution. A nested procedure is used, employing Newton-Krylov methods [11]. In the innermost loop, a linearized system is solved in the pre-conditioned Krylov-subspace. The outer loop uses a Newton-Raphson iteration to resolve the non-linearities. Besides being an effective general procedure for sparse systems, our implementation of parallel pre-conditioners also insures that the solver can run effectively on large-scale parallel computing architectures—the chief reason our code exhibits its high degree of scalability.

### 2.3. Pre-collapse Progenitor Models and Initial Conditions

For a progenitor model, we employ the widely used Woosley and Weaver [12] S15S7b2  $15M_\odot$  progenitor. The initial model is collapsed using RH1D, our 1-D Newtonian Lagrangean radiation hydrodynamics code. Immediately prior to core bounce, the iron core, the silicon shell, and a portion of the oxygen shell are zoned into a 256 radial-mass-zone mesh with zoning that is tuned to yield a high spatial resolution grid in the proto-neutron star and the inner 200 km of the collapsed core. This zoning sets up a radial grid that is compatible with subsequent 2-D Eulerian simulations. In both the 1-D collapse and the subsequent 2-D evolution with V2D, the neutrino-energy spectrum (range, 0–375 MeV) is discretized into 20 energy groups with group widths increasing geometrically with energy to resolve the Fermi surface of the electrons and neutrinos in the developing proto-neutron star. The initial values for  $T$ ,  $\rho$ , and  $Y_e$  are interpolated from the original S15S7B2 data onto the Lagrangean mass grid, and the initial radial coordinates of each mass shell are computed consistently with density.

### 2.4. Nuclear Microphysics and the Equation of State

Stellar-core collapse simulations require an equation of state (EOS) that handles a density range of roughly  $10^5$ – $10^{15}$  g cm $^{-3}$ , a temperature range of 0.1–25 MeV, and an electron-fraction range of 0.0–0.5. The EOS also must handle different regimes of equilibrium states. Throughout most of the core, the material is in nuclear statistical equilibrium (NSE), though matter in the silicon shell and beyond does not attain NSE until the bounce-shock wave passes through it.

For supernova simulations, we use the Lattimer-Swesty EOS [4,3] in tabular form. The thermodynamic quantities are tabulated in terms of independent variables: density,  $\rho$ ; temperature,  $T$ ; and electron fraction,  $Y_e$ . We have tabulated this EOS in a thermodynamically consistent way according to the prescription in Swesty [10]. We refer to this combination collectively as LS-TCT. Although the Lattimer-Swesty EOS is commonly used, and tabulations of it are also common, most tabular interpolations are not constructed to *guarantee* thermodynamic consistency. When a non-thermodynamically-consistent scheme is used, spurious entropy can be generated or lost.

### 2.5. Neutrino Microphysics

The energetics of a core-collapse supernova are dominated by neutrinos and, therefore, it is necessary to have accurate opacities and rates for the various important neutrino processes.

The collection of reactions that are important, or possibly important, to the supernova problem is large and evolving. Of undisputed importance is electron capture by protons and protons bound in nuclei, neutrino production through electron-positron pair annihilation, and non-conservative neutrino-electron scattering. The major contributors to neutrino opacity are coherent scattering of neutrinos from nuclei and conservative scattering from free nucleons.

The models presented here contain neutrino microphysics as described in Bruenn [1], with two exceptions. The neutrino-nucleus scattering opacity has been modified to take into account the form factor introduced by Burrows, Mazurek, and Lattimer [2], and we have neglected the effect of neutrino-antineutrino annihilation. Full neutrino-electron scattering, as presently described, is included in the code but is not turned on in the models described here. Additional effects, including nucleon recoil, ion-ion correlations, and flavor-changing reactions, are under active development as additions to these baseline models.

### 3. Parallel Implementation

The numerical solution of the radiation-hydrodynamics equations delineated in the prior sections is accomplished via the V2D code previously mentioned. The size of the problem we solve necessitates massively parallel computing resources. Since we solve a long-timescale problem, it is necessary that we achieve strong-scaling, *i.e.*, we wish the fixed-size problem to scale well to large numbers of processors to reduce our wall-clock time to solution. Also, the number of variables in the problem requires a large amount of memory. Parallelism in V2D is achieved via spatial domain decomposition of the 2-D spatial domain into a logically Cartesian topology of 2-D subdomains. Under this decomposition, our finite-difference algorithms for the hydrodynamic and radiation-transport equations require a limited set of communication patterns. The evaluation of the finite-difference stencils for both sets of equations requires only nearest-neighbor message-passing with this decomposition. Unfortunately, global reduction operations are needed in a few instances. The calculation of the timestep size  $\Delta t$  requires a global reduction to determine the minimum CFL time for the entire domain. The iterative solution, by Newton-Krylov methods, of the implicitly-differenced diffusion equations requires global reductions to evaluate vector inner products. In the BiCGSTAB algorithm, which is often used to solve the linear systems in Newton-Krylov methods, this can mean multiple global reductions per iteration. This can impose a bottleneck to scalability when carrying out simulations on large numbers of processors. To reduce this bottleneck, we have developed a restructured variant of the BiCGSTAB algorithm that is algebraically equivalent, but requiring only a single global reduction per iteration, and which also allows for better loop optimization. The effect of this improvement can be seen in Fig. 1, where we plot the parallel speedup of V2D on a supernova model on *seaborg*, the IBM-SP at NERSC.

The major floating point cost of the Newton-Krylov algorithm is expended in the multiplication of the Jacobian operator for the non-linear diffusion equation with a vector. This operation requires only nearest-neighbor communication to evaluate the finite-difference stencil of the divergence operator. We amortize this cost by performing the nearest neighbor-communication asynchronously while carrying out the portions of the matrix-vector multiply operation corresponding to the interior zones of each subdomain. This yields a further improvement in scalability as seen in Fig. 1.

V2D's data input and output is accomplished in parallel using the parallel HDF5 libraries, which in turn rely on the MPI-I/O interface. The entire simulation does parallel collective reads and writes into a single HDF5 file. This strategy ensures that checkpoint and restart files are independent of processor count. I/O is typically costly, and, therefore, we checkpoint relatively infrequently. We have also found that parallel I/O can become an impediment to scaling beyond 1024 processors, and




Figure 1. Parallel speedup of V2D on the NERSC IBM-SP (seaborg).

we are investigating strategies to mitigate this problem.

#### 4. A Few Scientific Discoveries...

Using V2D, we have carried out a simulations of the convectively unstable post-bounce epoch of a core-collapsed  $15M_{\odot}$  star. As mentioned previously, models were evolved through the collapse epoch using RH1D. Immediately prior to core bounce, when the central density is about  $10^{14} \text{ g cm}^{-3}$ , remapping to 2-D spherical polar coordinates takes place, keeping radial zoning unaltered. For the S15S7b model and  $K = 180 \text{ MeV}$  in the LS EOS, this corresponds to a time of 233.087 ms. The model is subsequently evolved further with V2D. This allows the shock to develop naturally on the Eulerian mesh and eliminates the possibility of spurious shocks in the quasi-static post-bounce core that have been observed when the 1D-to-2D transition is made at a later time.

An interesting discovery we have made is that by making the 1D-to-2D transition prior to bounce, fluid-instabilities develop earlier after the shock breakout stage. This is illustrated in Fig. 2, which shows a Rayleigh-Taylor instability developing in a region with a negative radial entropy gradient as a result of a weakening shock wave. This convective instability develops within 5–6 ms of




Figure 2. The development of Rayleigh-Taylor initiated convection in a region with negative entropy gradient. The color denotes the entropy per baryon while the streaks are a texture map that show the instantaneous structure of the velocity field.

the beginning of the simulation and core bounce. This is substantially earlier than reported by other groups who have started 2-D simulations after the prompt shock has stalled. The instability develops in the region where the material is optically thick to the bulk of the neutrinos. The convective circulation in this region causes a brief enhancement to the neutrino luminosity by advecting upward neutrinos that are “trapped” in the fluid on a diffusion timescale much longer than the fluid dynamic timescale. When the fluid containing the neutrinos reaches lower densities, they are able



to escape, since their opacity drops with density. This convective circulation also drags neutron rich material downward into the core as seen in Fig. 3, which shows the electron fraction of the material.




Figure 3. The electron fraction in the region below the neutrinosphere during the brief period of vigorous convection.

This period of vigorous circulation below neutrinosphere lasts only briefly and does not play a long term role in enhancing the luminosity. In fact, after 10 ms, the net flow of neutrinos in this region is downwards, since they are advected deeper into the core by high density matter that is quasi-statically settling after passing through the nearly-stalled accretion shock. This vigorous convection in the optically thick regions reduces the entropy gradient, after which convection in this region diminishes.

In this particular model, the stalled accretion shock does not revitalize on the timescale of this simulation (about 33 ms). Convection persists in the region above the neutrinosphere but below the stalled shock as is illustrated in Fig. 4. One can also see that by this time, the convection below the neutrinosphere has halted after it has stabilized the entropy gradient out to a radius of approximately 40 km. Whether there is sufficient reheating taking place in the convective region above the neutrinosphere to revitalize the shock in the long term, is the subject of our continuing investigations.




Figure 4. The entropy in the region below the neutrinosphere during the brief period of vigorous convection.

## 5. Conclusions and Future Directions

Through our simulations, we have found that the onset of fluid instabilities may occur earlier than previously thought. These instabilities occur as a result of the negative entropy gradient left in the wake of a weakening shock wave. We have also found that, in this model, vigorous protoneutron star convection does not persist for any significant time after the entropy gradient has been stabilized. Our results also indicate that shock revitalization by means of convective reheating enhancements does not occur on timescales of a few tens of milliseconds.

Our research into this mechanism is continuing at a brisk pace. We are moving beyond this baseline model to examine the effects of various microphysics changes within the model including the variation of the nuclear force parameters within the equation of state. Another area in which we are accelerating our research is in the development of 3-D models. It is unclear whether the convective structure we currently see in 2-D models will persist in 3-D. In addition, we hope to explore the effects of rotation with 3-D models. These studies will continue to occupy us beyond terascale computing and well into the petascale era.

## Acknowledgments

The authors thank numerous people who have aided our supernova modeling efforts. Foremost on this list are TSI team members Jim Lattimer, Ken DeNisco, Amy Irwin, Chris Tartamella, and

Clint Young (SUNY Stony Brook); Ed Bacht and Polly Baker (Indiana University at Indianapolis); Dennis Smolarski (Santa Clara University); John Blondin (North Carolina State University). We also thank members of some of the SciDAC funded ISIC and SAPP teams that have contributed technologies that have made our efforts easier: Micah Beck and Scott Atchley (LOCI Lab, University of Tennessee Knoxville); Xiaowen Xin and Terence Critchlow (SDM ISIC Team, LLNL); Dan Reynolds and Carol Woodward (TOPs ISIC Team, LLNL). Finally, we acknowledge the many members of the NERSC staff who have worked to help us carry out the computations described in this paper including Horst Simon, Francesca Verdier, Wes Bethel, David Skinner, Richard Gerber, John Shalf, Eli Dart, and Brent Draney.

We gratefully acknowledge the support of the U.S. Department of Energy, through SciDAC Award DE-FC02-01ER41185, by which this work was funded. We are also grateful to the National Energy Research Scientific Computing Center (NERSC) for computational support.

## References

1. Bruenn S W 1985 *Astrophys. J. Supp.*, **58**, 771
2. Burrows A, Mazurek T J and Lattimer J M 1981 *Astrophys. J.*, **251**, 325
3. Lattimer J M, Pethick C J, Ravenhall D G and Lamb D Q 1985 *Nucl. Phys. A*, **432**, 646
4. Lattimer J M and Swesty F D 1991 *Nucl. Phys. A*, **535**, 331
5. Levermore C D and Pomraning G C 1981 *Astrophys. J.*, **248**, 321
6. Myra E S, Bludman S A, Hoffman Y, Lichtenstadt I, Sack N and Van Riper K A 1987 *Astrophys. J.*, **318**, 744
7. Stone J M and Norman M L 1992 *Astrophys. J. Supp.*, **80**, 753
8. Stone J M and Norman M L 1992 *Astrophys. J. Supp.*, **80**, 791
9. Stone J M, Mihalas D and Norman M L 1992 *Astrophys. J. Supp.*, **80**, 819
10. Swesty F D 1996 *J. Comp. Phys.*, **127**, 118
11. Swesty F D, Smolarski D C and Saylor P E 2004 *Astrophys. J. Supp.*, **153**, 369
12. Woosley S E and Weaver T A 1995 *Astrophys. J. Supp.*, **101**, 181

